INFORMATION SYSTEMS EDUCATION JOURNAL

Volume 24, No. 4
July 2026
ISSN: 1545-679X

In this issue:

4. **Revisiting Academic Source Selection: Habit, Dependency, and Cognitive Effort in the Age of GenAI**
Muqing Guo, The University of Melbourne
Winn Chow, The University of Melbourne
Hasti Samadi, The University of Melbourne
Arzoo Atiq, The University of Melbourne
Tanya Linden, The University of Melbourne

22. **AI in IS Education: A Categorization of ISCAP Publications Using Zuboff's Automate–Informate–Transformate Framework**
Jason H. Sharp, Hardin-Simmons University
John E. Anderson, Utah Valley University
Guido Lang, Quinnipiac University

36. **Toward Curriculum-Aligned Conversational NPCs for Cybersecurity Training Simulations: A Feasibility Framework and Prompt-Based Comparative Evaluation**
Chukwuemeka Ihekweazu, Sam Houston State University
Naomi Aghado, Viterbo University
Iheme H. S. Ihekweazu, Nasarawa State University
Elizabeth Adepeju Adelowo, St. Cloud State University

54. **An Intersectional Approach to Exploring Computing Academic Success Self-Efficacy (CASS) for Black Undergraduate Students**
Sonal Dekhane, Georgia State University
Nannette P. Napier, Georgia State University

67. ***Teaching Case***
Bridging the Gap via Community Engagement: Enhancing Systems Analysis & Design Education through Client-Based Projects
Reda Haddouch, University of Montana
Clayton Looney, University of Montana
Eric Legvold, United Way of Missoula County
Shawn Clouse, University of Montana
Fang Chen, University of Montana

97. **An Exploratory Analysis of Perceptions and Uses of Generative AI Tools for Higher Education**
Karen Paullet, Robert Morris University
Adnan Chawdhry, PennWest University
Jamie Pinchot, Robert Morris University

iscap
Information Systems & Computing Academic Professionals

The **Information Systems Education Journal** (ISEDJ) is a double-blind peer-reviewed academic journal published by **ISCAP** (Information Systems and Computing Academic Professionals). Publishing frequency is five times per year. The first year of publication was 2003.

ISEDJ is published online (https://isedj.org). Our sister publication, the Proceedings of the ISCAP Conference (https://iscap.us/proceedings) features all papers, abstracts, panels, workshops, and presentations from the conference.

The journal acceptance review process involves a minimum of three double-blind peer reviews, where both the reviewer is not aware of the identities of the authors and the authors are not aware of the identities of the reviewers. The initial reviews happen before the ISCAP conference. All papers, whether award-winners or not, are invited to resubmit for journal consideration after applying feedback from the Conference presentation. Award winning papers are assured of a publication slot; however, all re-submitted papers including award winners are subjected to a second round of three blind peer reviews to improve quality and make final accept/reject decisions. Those papers that are deemed of sufficient quality are accepted for publication in the ISEDJ journal. Currently the target acceptance rate for the journal is under 35%.

Information Systems Education Journal is pleased to be listed in the Cabell's Directory of Publishing Opportunities in Educational Technology and Library Science, in both the electronic and printed editions. Questions should be addressed to the editor at editor@isedj.org or the publisher at publisher@isedj.org. Special thanks to volunteer members of ISCAP who perform the editorial and review processes for ISEDJ.

2026 ISCAP Board of Directors

Amy Connolly James Madison University President	Michael Smith Georgia Institute of Technology Vice President	Jeff Cummings Univ of NC Wilmington Past President
David Firth University of Montana Director	Mark Frydenberg Bentley University Director/Secretary	Leigh Mutchler James Madison University Director
RJ Podeschi Millikin University Director/Treasurer	Bryan Reinicke Rochester Institute of Technology / Director	Jeffry Babb West Texas A&M University Director/Curricular Matters
Eric Breimer Siena University Director/2026 Conf Chair	Thomas Janicki Univ of NC Wilmington Director/Meeting Planner	Xihui "Paul" Zhang University of North Alabama Director/JISE Editor

Copyright © 2026 by Information Systems and Computing Academic Professionals (ISCAP). Permission to make digital or hard copies of all or part of this journal for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial use. All copies must bear this notice and full citation. Permission from the Editor is required to post to servers, redistribute to lists, or utilize in a for-profit or commercial use. Permission requests should be sent to Kevin Mentzer, Editor, editor@isedj.org.

INFORMATION SYSTEMS EDUCATION JOURNAL

Editors

Kevin Mentzer
Editor
Nichols College

Ira Goldman
Associate Editor
Siena University

David Yates
Associate Editor
Bentley University

Michelle Louch
Teaching Cases & Exercises
Editor
University of Pittsburgh - Greensburg

Mark Pisano
Teaching Cases & Exercises
Associate Editor
Southern Connecticut
State University

Thomas Janicki
Publisher
Univ of NC Wilmington

David Woods
Assistant Publisher
Miami University
Regionals

Paul Witman
Emeritus Editor
(2021-2026)
California Lutheran
University

Jeffry Babb
Emeritus Editor
(2016-2021)
West Texas A&M
University

Donald Colton
Emeritus Editor
(2003-2010)
Brigham Young University
Hawaii

Toward Curriculum-Aligned Conversational NPCs for Cybersecurity Training Simulations: A Feasibility Framework and Prompt-Based Comparative Evaluation

Chukwuemeka Ihekweazu
cei004@shsu.edu
Computer Science Department
Sam Houston State University
Huntsville, TX 77340, USA

Naomi Aghado
naomiaghado@yahoo.com
College of Nursing and Health Department
Viterbo University
La Crosse, WI 54601, USA

Iheme H.S Ihekweazu
eihekweazu@gmail.com
Department of Information Technology
Nasarawa State University
Nasarawa, Nigeria

Elizabeth Adepeju Adelowo
eaadelowo@stcloudstate.edu
Educational Leadership and Higher Education Department
St. Cloud State University
St Cloud, MN 56301, USA

Abstract

Simulation-based and gamified platforms have become central to cybersecurity education; however, numerous training environments remain procedurally static and offer limited real-time instructional scaffolding. This constraint is particularly significant in self-paced learning contexts, where novice learners require adaptive guidance to navigate decision-rich tasks such as phishing analysis, log interpretation, and incident response triage. This study proposes a feasibility framework for integrating AI-driven Non-Playable Characters (NPCs) as conversational guides within cybersecurity training simulations. The framework combines a curriculum-informed knowledge base aligned with established workforce and threat frameworks such as NIST NICE and MITRE ATT&CK, a large language model (LLM)-enabled conversational layer for real-time dialogue management, and a simulation interface deployable across browser-based and game-engine environments. To evaluate feasibility and instructional alignment, the study conducts a prompt-based comparative assessment of two NPC designs: a curriculum-informed agent grounded in structured course materials and a general-purpose LLM agent without domain-specific tuning. The results indicate that the curriculum-informed NPC generates more instructionally accurate, standards-aligned, and pedagogically structured responses across representative cybersecurity scenarios. This work offers a scalable design model and evaluation methodology that support future empirical research and the development of immersive, AI-guided cybersecurity training systems.

Keywords: cybersecurity education; conversational AI; non-playable characters; simulation-based training; instructional scaffolding; large language models

Recommended Citation: Ihekweazu, C., Aghado, N., Ihekweazu, I., Adelowo, E., (2026). Toward Curriculum-Aligned Conversational NPCs for Cybersecurity Training Simulations: A Feasibility Framework and Prompt-Based Comparative Evaluation. *Information Systems Education Journal* v24(n4) pp 36-53. DOI# https://doi.org/10.62273/KRUV1548

Toward Curriculum-Aligned Conversational NPCs for Cybersecurity Training Simulations: A Feasibility Framework and Prompt-Based Comparative Evaluation

Chukwuemeka Ihekweazu, Naomi Aghado, Iheme H.S Ihekweazu and Elizabeth Adepeju Adelowo

1. INTRODUCTION

This journal version extends the conference proceedings by introducing a structured rubric-based evaluation approach, expanded validity analysis, and enhanced standards-aligned design implications for curriculum-informed conversational NPCs.

Cybersecurity education increasingly depends on simulation-based learning environments to help learners develop practical skills in threat detection, incident response, and defensive analysis without the risks associated with live systems. Virtual labs, cyber ranges, and gamified platforms offer structured opportunities for experiential learning and deliberate practice. However, despite widespread adoption, many cybersecurity training environments remain procedurally static and provide limited real-time instructional support, especially in self-paced settings where instructor availability is restricted. Consequently, learners may struggle to interpret artifacts, choose appropriate investigative actions, or understand why a particular decision is correct, which can slow skill development and decrease motivation.

A persistent challenge in simulation-based cybersecurity instruction is the lack of adaptive, context-sensitive feedback mechanisms that respond to learner intent and decision-making in real time. In many platforms, guidance is delivered through fixed walkthroughs, scripted hints, or external documentation. While these resources can support task completion, they often fail to provide individualized scaffolding that aligns with a learner's current understanding, cognitive load, or error patterns. Educational theory suggests that learners benefit from timely feedback, structured progression, and reflective prompts that promote knowledge construction and transfer. In cybersecurity, where tasks are dynamic and decision-rich, the absence of responsive instructional support can lead to frustration and surface-level completion rather than deep comprehension.

Recent advances in conversational artificial intelligence, particularly large language models (LLMs), have enabled interactive dialogue systems capable of producing natural language explanations, answering questions, and adapting responses based on conversational context. In educational settings, conversational agents have demonstrated potential to support engagement and self-directed learning when designed with pedagogical intent. Yet, most implementations remain external to immersive simulation environments, operating as standalone chatbots rather than integrated instructional actors. This separation limits their capacity to deliver situationally grounded guidance within realistic cyber scenarios where learners must interpret evidence, apply frameworks, and make sequential decisions.

To address this gap, this study proposes the use of AI-driven Non-Playable Characters (NPCs) as embedded conversational guides within cybersecurity training simulations. Unlike static NPCs or pre-scripted characters, AI-driven NPCs can adopt role-based instructional behaviors (e.g., mentor, adversary, incident handler) and respond dynamically to learner actions using dialogue grounded in curriculum objectives and recognized cybersecurity frameworks. The proposed feasibility framework integrates structured domain knowledge, LLM-enabled dialogue management, and a simulation interface capable of deployment in web-based and game-engine environments. The framework is designed to support scalable instructional scaffolding while maintaining alignment with workforce competencies and adversarial tactics.

This paper presents a feasibility-oriented design and evaluation approach for conversational NPCs in cybersecurity education. Rather than measuring learner outcomes through human-subject experimentation, the study conducts a prompt-based comparative evaluation of two NPC designs: (1) a curriculum-informed NPC grounded in structured educational materials and aligned with NIST NICE and MITRE ATT&CK, and (2) a general-purpose LLM NPC without domain-specific tuning. By comparing instructional accuracy, clarity, scaffolding quality, and standards alignment across representative scenarios, this work provides foundational

©2026 ISCAP (Information Systems and Computing Academic Professionals)

evidence supporting the viability of curriculum-aligned conversational NPCs as scalable instructional components in cybersecurity simulations.

2. PURPOSE AND CONTRIBUTIONS

Purpose: This study proposes a feasibility framework for integrating AI-driven non-playable characters (NPCs) as conversational guides within cybersecurity training simulations, with an emphasis on curriculum alignment, instructional scaffolding, and standards-informed interaction design.

Contributions. This paper contributes: (1) a modular architecture for embedding conversational NPCs into cybersecurity simulations across web-based and game-engine platforms; (2) a curriculum-informed NPC design approach grounded in structured instructional materials and aligned with NIST NICE and MITRE ATT&CK; and (3) a structured, rubric-driven prompt-based comparative evaluation method that enables transparent and reproducible assessment of instructional accuracy, scaffolding quality, and standards alignment between curriculum-informed and general-purpose LLM agents.

3. RESEARCH QUESTIONS

The research is guided by the following questions:

1. How can AI-driven NPCs simulate realistic and pedagogically meaningful interactions within cybersecurity training environments?
2. What instructional roles can conversational NPCs perform to support learning across exploration, practice, and reflection?
3. How can instructional design principles be embedded within AI NPCs to enhance engagement and personalization?

4. SCOPE OF STUDY

This study focuses on the development of a framework architecture, sample implementation scenarios, and pedagogical justification grounded in existing literature. The outcome is a proposed model that can serve as a foundation for future system development and empirical studies in cybersecurity education.

5. LITERATURE

Conversational AI in Education and Training Environments

Recent advancements in natural language processing (NLP) and large language models (LLMs) have enabled the development of conversational agents capable of interacting with learners in natural, context-aware ways. These agents commonly deployed as chatbots or virtual tutors have shown promise in supporting self-directed learning, answering questions, and providing timely feedback (Fryer & Nakao, 2020; Winkler & Söllner, 2018). When conversational agents are designed with pedagogical intent, they can enhance engagement and foster reflective learning in STEM education (Kim, 2019). However, current applications tend to exist outside immersive environments, limiting their potential to support real-world learning contexts like cybersecurity.

In cybersecurity, where learners must acquire both procedural knowledge and situational awareness, static learning resources fall short. Intelligent agents must simulate pedagogically meaningful interactions through adaptive dialogue, in which NPCs provide guided prompts and context-sensitive responses tailored to learners' decisions, thereby supporting cognitive scaffolding (VanLehn, 2011; Vygotsky, 1978). The shift toward AI NPCs embedded in simulations offers a new opportunity for real-time, contextualized guidance that responds to learner decisions and evolving threat scenarios.

In their paper, Mierzwa et al. (2019) demonstrated that feasibility studies can introduce AI chatbots as viable tools before participant testing, highlighting their potential to improve efficiency and support user interaction. Following this precedent, the current NPC framework should likewise be regarded as a proof-of-concept, illustrating technical and instructional feasibility ahead of large-scale validation.

While conversational AI has seen increased use in general education, a review by Ahmed and Parsons (2021) found that most cybersecurity simulations are static, relying on hard-coded instruction paths and offering limited interactivity. Similarly, Yin, Zhu, and Wang (2022) concluded that adaptive instruction and intelligent agents remain underutilized in cybersecurity education research. These environments fail to reflect the dynamic, decision-rich nature of real-world cyber defense,

highlighting an urgent gap for more context-aware instructional tools.

Non-Playable Characters (NPCs) in Simulations and Game-Based Learning

NPCs have long played important roles in game-based learning by acting as guides, mentors, adversaries, or story-driving characters. In educational contexts, NPCs can help learners practice critical thinking, apply problem-solving strategies, and experience failure in low-risk environments (Gee, 2003; Dede, 2009). However, in cybersecurity education, NPCs are often limited to non-interactive roles or pre-scripted interactions that fail to reflect real-world complexity.

Emerging research in immersive learning environments supports the integration of intelligent NPCs that adapt based on learner behavior and simulate realistic cyber threats (Moser et al., 2021). By assigning NPCs instructional roles across phases of learning exploration, practice, and reflection, they can serve as scaffolds that activate prior knowledge, offer formative feedback, and encourage self-regulation (Kolb, 1984; Herrington & Oliver, 2000).

Platforms such as TryHackMe and RangeForce offer valuable scenario-based exercises. However, unlike the proposed framework, they lack embedded conversational NPCs capable of adaptive, role-based instruction. This highlights the novelty of embedding AI-driven NPCs into simulation-based cyber training. Learners must rely on written walkthroughs, forums, or static hints, which significantly limit immediacy and personalization (Dodge, Ragsdale, & Reynolds, 2020; Yamin, Katt, & Gkioulos, 2020). This design may be suitable for self-paced learners with a technical background, but it provides limited instructional support for novices or those who require adaptive guidance.

In this study, the authors analyzed survey data from 1,597 first-year students across multiple semesters, revealing a rapid increase in the adoption of generative AI tools, such as ChatGPT, for academic tasks. Students primarily used AI for homework and research preparation. The study highlights the importance of data literacy in ensuring equitable and ethical adoption of AI. These findings underscore the rapid penetration of AI-driven tools into educational spaces, underscoring the timeliness of integrating NPCs as conversational AI assistants in cybersecurity education (Frydenberg, Mentzer, & Patterson, 2026).

Instructional Design Principles Embedded in AI NPCs

Embedding instructional design principles into AI NPCs involves more than enabling them to speak or respond. It requires that their behavior, dialogue, and interactions follow pedagogically intentional models. Models like Gagné's Nine Events of Instruction (Gagné et al., 2005), Merrill's First Principles (Merrill, 2002), and the ARCS Motivation Model (Keller, 2009) provide frameworks for designing learning sequences that capture attention, build relevance, support mastery, and provide feedback.

For example, an NPC acting as a mentor may follow a pattern of instruction by: (1) gaining attention with a compelling opening prompt, (2) informing the learner of objectives, (3) guiding learning with tips or mini-tasks, (4) providing feedback based on learner action, and (5) encouraging reflection at scenario end.

From a personalization standpoint, AI NPCs can leverage learner modeling, where past decisions, task success rates, and interaction history are analyzed to customize instructional content. However, adaptive instructional behavior must remain transparent and interpretable to learners and instructors to support trust, accountability, and instructional integrity. Explainable AI (XAI) research in educational contexts emphasizes that AI-supported learning systems should provide understandable reasoning for their outputs and decisions, particularly when personalization influences learner outcomes (Khosravi et al., 2022). This is particularly relevant in cybersecurity training, where inaccurate or non-transparent guidance may unintentionally reinforce unsafe investigative practices. By embedding Universal Design for Learning (UDL) principles such as offering multiple means of representation and engagement, NPCs can deliver information in varied formats (textual, visual, interactive) and adjust pacing or challenge level based on learner performance (CAST, 2018).

Engagement is further strengthened through affective adaptation, where AI systems monitor and respond to learner emotions using methods from affective computing, such as sentiment analysis and interaction timing. Although empirical testing with participants was not conducted, frustration was conceptualized from prior literature as observable through long response delays, repeated errors, or disengagement. For example, an NPC might adapt by shifting from a high-pressure dialogue

style to a more supportive tone. Serholt et al. (2021) emphasize that emotionally aware AI agents can improve persistence and satisfaction in simulation-based learning environments.

Jiang and Nakatani (2025) present an empirical study integrating Generative AI tools (e.g., ChatGPT, Gemini) into IS coursework. Their results show that students using GenAI performed as well as or better than those in control groups. Importantly, concerns of academic dishonesty and overreliance did not manifest significantly. Students responded favorably to AI-assisted assignments, reinforcing the feasibility of conversational AI in education.

Most existing literature on educational chatbots and AI tutors focuses on general-purpose systems for answering questions or guiding students through structured content, Chassignol et al. (2018). However, there is limited research on NPCs that assume distinct instructional roles, such as adversaries simulating phishing attacks, mentors providing reflective feedback, or incident handlers guiding triage decisions. This type of pedagogically aligned, character-based simulation is especially relevant to cybersecurity but remains largely unexplored in both research and practice.

While prior research highlights the benefits of conversational agents and NPCs in education, few studies integrate them with instructional design models tailored for cybersecurity simulations. There is a significant opportunity to develop AI-driven NPCs that combine situational fidelity with instructional purpose providing personalized, scaffolded, and meaningful interactions throughout the learning experience. In summary, while prior research has explored conversational AI and NPCs in education separately, their integration in cybersecurity training particularly through differentiated instructional roles embedded in real-time simulations remains underdeveloped. This study addresses this critical gap by proposing a unified framework that brings together instructional design principles, conversational AI capabilities, and simulation-based learning in a novel and scalable way.

6. METHODOLOGY

This study introduces a conceptual framework in Figure 1 (in the appendices) for integrating conversational AI-driven non-playable characters (NPCs) into cybersecurity simulations, with a specific focus on instructional design. The framework is grounded in the principles of scaffolding, adaptive feedback, and contextualized learning, and is designed to be applicable across web-based or immersive platforms.

At its core, the framework consists of two role-differentiated AI agents developed using the Convai platform. Development required approximately 25 hours, including data preparation, persona scripting, and integration. Implementation also required familiarity with Unity or Unreal Engine for simulation integration and character deployment. Future researchers can expand upon this work by importing new datasets, building more sophisticated conversational algorithms, and customizing character roles:

- **NPC 1 (Curriculum-Informed Agent):** This agent, in Figure 2, is built using structured, locally sourced educational content such as syllabi, lab instructions, and NIST (2020) NICE-aligned course materials as seen in its "Knowledge Bank" in Figure 4. Its function is to offer highly specific, standards-aligned instructional responses. By anchoring this agent in pre-defined learning outcomes, it can deliver accurate and consistent guidance to learners in scenario-based tasks.

- **NPC 2 (General LLM-Based Agent):** This agent, in Figure 3, operates using a generalized large language model without local customization. While capable of responding broadly to cybersecurity-related prompts, its answers are not optimized for curriculum alignment or specific instructional design goals. It represents a baseline for comparison and reflects how most open-ended AI tools function in educational contexts.

These agents are embedded within a three-layer architecture as seen in Figure 1:

1. **Input Layer:** Defines the foundational knowledge base for each NPC. For NPC 1, this includes structured datasets derived from vetted course content, lab instructions, and standards-aligned materials (e.g., NIST NICE, MITRE ATT&CK). For NPC 2, this layer consists of generalized LLM knowledge without domain-specific constraints.

Scenario Context Injection (Dynamic Sub-layer):
In addition to the static knowledge base, each interaction session incorporates a *scenario-*

specific context payload that dynamically constrains NPC responses. This context includes:

- Task objectives (e.g., perform port scan and validate HTTP service)
- Known environmental artifacts (e.g., open ports, IP addresses, logs)
- Expected learning outcomes and constraints

This context is injected at runtime through structured prompts and session memory, ensuring that NPC responses remain grounded in the active lab scenario rather than general training data. This mechanism enables both NPC configurations to respond specifically to a given exercise while preserving their respective instructional characteristics.

2. **Dialogue Management Layer:** Uses Convai's session-based interaction engine to interpret learner prompts and generate real-time NPC dialogue. This layer manages conversational flow, persona behavior, response timing, and contextual continuity across interactions.

3. **Instructional Output Layer:** Transmits responses through a simulation interface (web-based or immersive environment). NPCs assume defined instructional roles (e.g., mentor, analyst, adversary) and provide context-sensitive guidance aligned with learner actions.

This modular architecture enables the comparison of NPC behavior in response to the same scenario prompts. These prompts are aligned with cybersecurity competencies such as phishing detection, network log analysis, and incident response triage. A controlled input-output comparison reveals how each NPC supports learners across dimensions of instructional clarity, engagement, and fidelity to established standards (e.g., NIST (2020) NICE, MITRE (n.d.) ATT&CK). By combining pedagogical intent with real-time dialogue generation, the framework offers a scalable model for embedding intelligent agents into cybersecurity training environments. It also serves as a proof of concept for applying AI-enhanced scaffolding to complex, domain-specific simulations.

7. EVALUATION AND VALIDATION

This study adopts a feasibility-oriented, prompt-based comparative evaluation approach to assess the instructional viability of AI-driven non-playable characters (NPCs) as conversational guides in cybersecurity training simulations. Because the present work is conceptual and does not involve human subjects, evaluation is conducted through controlled scenario prompts designed to simulate common learner interactions in cybersecurity training contexts.

Two NPC configurations are evaluated:

a. NPC 1 (Curriculum-Informed Agent): A structured agent grounded in locally curated educational materials (e.g., lab instructions, course content, and standards-aligned learning objectives).

b. NPC 2 (General LLM Agent): A general-purpose conversational agent powered by a large language model (LLM) without curriculum-specific grounding.

Both NPCs were presented with identical prompts representing core cybersecurity learning tasks (e.g., phishing identification, incident response reasoning, and Windows log interpretation). Responses were captured and reviewed side-by-side to examine differences in instructional accuracy, alignment with standards, and pedagogical scaffolding.

Evaluation Design (Prompt-Based Comparative Method)

The evaluation reflects realistic instructional interactions that occur during simulation-based training. Prompts were intentionally designed to require reasoning, explanation, and procedural guidance rather than simple factual recall. This design supports analysis of instructional depth, clarity, and scaffolding.

Evaluation Governance and Scoring Procedure

To improve transparency and reduce potential bias, NPC responses were anonymized during the initial evaluation to mitigate confirmation bias. Responses were reviewed side-by-side using a structured rubric. Scoring was conducted using predefined rubric criteria applied consistently across all scenarios. Evaluation was conducted by a single domain-informed evaluator with experience in cybersecurity instruction and digital forensics. This approach is consistent with feasibility-oriented studies, in which the objective is a structured comparison rather than statistical generalization. Scoring decisions were guided by explicit rubric definitions to ensure consistency and reduce subjective variation. To support reproducibility, all prompts, NPC configurations, and scoring criteria are provided in the appendices. Future work will incorporate multiple evaluators and inter-rater reliability measures

(e.g., Cohen's kappa) to further strengthen scoring reliability.

Evaluation Rubric

To increase rigor and reproducibility, NPC responses were evaluated using a structured rubric consisting of five criteria. Each criterion was scored on a 0–2 ordinal scale (0 = not demonstrated, 1 = partial, 2 = clearly demonstrated), where higher values indicate stronger instructional performance.

The rubric criteria are defined as follows:

a. Instructional Accuracy: Correctness and technical appropriateness of the response for the given cybersecurity task.

b. Standards / Framework Alignment: Evidence that the response reflects recognized cybersecurity frameworks (e.g., NIST NICE, MITRE ATT&CK) or competency-aligned guidance.

c. Scaffolding & Guidance Quality: Presence of stepwise support, formative prompting, and progression-based guidance suitable for learner development.

d. Clarity & Coherence: Logical structure, readability, and actionable communication.

e. Role Fidelity: Consistency in maintaining an instructional role (e.g., mentor, analyst, incident responder) and delivering role-appropriate guidance.

This rubric supports a transparent comparison between the curriculum-informed and general-purpose NPC designs by operationalizing instructional quality into measurable categories suitable for feasibility evaluation.

8. RESULTS

The comparative evaluation indicates that the curriculum-informed conversational NPC (NPC 1) demonstrates more consistent instructional alignment, structured scaffolding, and standards-informed reasoning than the general-purpose LLM NPC (NPC 2) across representative cybersecurity training prompts. Overall, NPC 1 demonstrated higher instructional accuracy, clearer scaffolding, and stronger alignment with standards-informed cybersecurity reasoning. NPC 2 produced generally fluent and relevant responses; however, it more frequently lacked procedural specificity, instructional sequencing, and explicit alignment with cybersecurity frameworks.

Table 1 summarizes the rubric-based comparative scoring approach and illustrates the pattern of performance differences observed across scenarios. Scores reflect rubric-based qualitative coding intended to support comparative feasibility analysis rather than statistical measurement.

NPC 1 achieved consistently higher rubric totals across all scenario prompts, with the largest differences observed in scaffolding quality and standards alignment

Across all evaluated scenarios, NPC 1 exhibited stronger consistency in instructional structure, often providing stepwise guidance and context-sensitive explanations aligned with cybersecurity competencies. For example, in phishing identification tasks, NPC 1 emphasized observable indicators and framed responses as repeatable workflows. NPC 2, while correct in many cases, tended to provide generalized explanations without structured instructional progression.

Similarly, for incident response reasoning prompts, NPC 1 delivered structured responses that reflected the expected phases of incident handling and included actionable recommendations appropriate for simulation-based training contexts. NPC 2 provided broadly relevant content but was more likely to omit key steps or present concepts without instructional order, reducing its utility as a scaffolded guide for learners.

Because this study is feasibility-oriented and uses a limited set of structured prompts, results are presented as descriptive comparisons rather than statistical inference. The consistent differences observed across scenarios suggest a clear pattern of stronger instructional alignment for the curriculum-informed NPC under controlled conditions. These findings support the feasibility of curriculum-aligned conversational NPCs as instructional components in simulation-based cybersecurity training environments.

9. LIMITATIONS

This study presents a conceptual framework for integrating AI-driven non-playable characters (NPCs) into cybersecurity training simulations, and as such, its generalizability is limited. The research does not involve empirical testing with

learners; instead, it relies on comparative evaluations of two NPC types (curriculum-informed vs. generalized LLM-based). As such, the findings illustrate feasibility rather than demonstrating learner outcomes, leaving aspects such as engagement, usability, and knowledge retention for future validation and evaluation.

Additionally, the framework was limited to two agent designs and developed on the Convai platform, introducing platform-specific constraints and limiting generalizability. While aligned with cybersecurity standards such as NIST (2020) NICE, and MITRE (n.d.) ATT&CK, the framework has not yet been extended to specialized domains (e.g., healthcare or finance), nor does it address ethical considerations like transparency, bias, or data privacy. These limitations suggest that the current work should be regarded as a proof-of-concept to guide future empirical testing, cross-domain adaptation, and the integration of ethical safeguards.

10. ETHICAL AND INSTRUCTIONAL INTEGRITY CONSIDERATIONS

Conversational AI agents deployed in cybersecurity training environments introduce instructional integrity risks, including hallucinated technical guidance, overconfident explanations, and inconsistent standards alignment. To mitigate these risks, curriculum-informed NPC designs should prioritize grounded knowledge sources, constrained response policies, and role-specific prompting that emphasizes safe, standards-informed procedures. Additionally, simulation platforms should transparently disclose AI assistance to learners, avoid the collection of unnecessary personal data, and implement logging controls that align with privacy and institutional requirements. These safeguards are essential to maintain learner trust and ensure that AI-guided instruction supports, rather than undermines, cybersecurity training outcomes.

11. THREATS TO VALIDITY

This study proposes and evaluates a feasibility framework for integrating AI-driven non-playable characters (NPCs) as conversational guides in cybersecurity training simulations using a prompt-based comparative method. While the approach provides a structured basis for evaluating instructional viability without human subjects, several threats to validity should be considered when interpreting the findings.

Internal Validity
The comparative evaluation relies on predefined scenario prompts that simulate learner interactions. As a result, the phrasing, complexity, and specificity of prompts may influence response quality and potentially advantage one NPC configuration over the other. To reduce this risk, identical prompts were issued to both NPCs and were selected to reflect common cybersecurity learning tasks (e.g., phishing identification, incident response reasoning, and log interpretation). However, prompt design remains a source of potential bias and may not fully capture the variability of authentic learner questioning during real training sessions.

Construct Validity
The evaluation rubric operationalizes instructional quality using five criteria: instructional accuracy, standards alignment, scaffolding quality, clarity, and role fidelity. Although these criteria are grounded in established instructional design principles and cybersecurity education needs, they do not directly measure learner outcomes such as knowledge retention, skill transfer, or long-term performance improvements. Additionally, constructs such as scaffolding and role fidelity require qualitative interpretation and may not fully capture the complexity of instructional effectiveness across diverse classroom and cyber range settings.

External Validity
Generalizability is limited by the scope of scenarios and the specific implementation context. The evaluation focuses on representative prompts rather than a comprehensive set of cybersecurity competencies, and results may not extend to specialized domains such as malware analysis, cloud forensics, OT/ICS security, or advanced threat hunting. Furthermore, platform-specific implementation constraints (e.g., differences across simulation engines or conversational frameworks) may influence deployment feasibility and response behavior. Because the study does not include human participants, the findings should be interpreted as feasibility evidence rather than validated learning impact.

Reliability
Reliability may be influenced by subjectivity in rubric-based scoring, particularly for criteria requiring judgment (e.g., scaffolding quality and role fidelity). If scoring is conducted by a single evaluator, results may reflect individual interpretation rather than consistent measurement. Future work should incorporate multiple evaluators, inter-rater agreement procedures, and larger prompt sets to strengthen

scoring consistency and improve reproducibility.

Future work will address these validity considerations through controlled learner studies, expanded scenario coverage, and multi-rater rubric scoring.

12. CONCLUSIONS

This study proposed a pedagogically grounded framework for integrating conversational AI-driven non-playable characters (NPCs) into cybersecurity training simulations. By designing and comparatively evaluating two distinct NPCs, one informed by structured curriculum data and the other based on a generalized large language model (LLM), the study highlights critical differences in instructional performance. The curriculum-aligned NPC demonstrated superior accuracy, scaffolding capabilities, and alignment with cybersecurity training standards, while the LLM-based NPC, though flexible, lacked instructional precision and relevance.

These findings underscore the importance of domain-specific instructional design when deploying AI agents in educational settings. Embedding structured pedagogical intent into conversational AI systems significantly enhances their ability to deliver relevant, standards-aligned learning support, particularly in complex, technical domains such as cybersecurity.

While this research is conceptual and does not involve human learners, the evaluation method establishes a foundation for future empirical studies. The framework can be expanded to include multimodal feedback (e.g., visual cues, interactive hints), adaptive learning pathways, and learner modeling to personalize instruction based on individual skill levels.

13. FUTURE WORK

Further research should explore the ethical implications of deploying AI NPCs in educational settings, including concerns about bias, transparency, and data privacy. As AI continues to transform digital learning environments, ensuring instructional integrity and learner trust will be critical in the development of intelligent, scalable educational tools. Additionally, future implementations may explore integration into virtual reality (VR) and augmented reality (AR) environments to further increase realism and engagement.

14. ACKNOWLEDGEMENTS

Generative AI tools (ChatGPT, OpenAI) were used to support language refinement and improve clarity during manuscript preparation. All technical content, interpretations, and final editorial decisions were reviewed and verified by the authors. The authors remain fully responsible for the accuracy, originality, and integrity of the work.

15. REFERENCES

Ahmed, R., & Parsons, D. (2021). Cybersecurity skills training: An analysis of online cybersecurity labs for educational use. *Education and Information Technologies, 26*(4), 4471–4494. https://doi.org/10.1007/s10639-021-10552-4

CAST. (2018). Universal Design for Learning Guidelines version 2.2. https://udlguidelines.cast.org/more/downloads/

Chassignol, M., Khoroshavin, A., Klimova, A., & Bilyatdinova, A. (2018). Artificial intelligence trends in education: A narrative overview. *Procedia Computer Science, 136*, 16–24. https://doi.org/10.1016/j.procs.2018.08.233

Dakpa, T., & Augustine, P. (2017). Study of phishing attacks and preventions. International Journal of Computer Applications, 163, 5–8. https://doi.org/10.5120/IJCA2017913461

Dede, C. (2009). Immersive interfaces for engagement and learning. *Science, 323*(5910), 66–69. https://doi.org/10.1126/science.1167311

Dodge, R. C., Ragsdale, D. J., & Reynolds, C. (2020). Cyber ranges: Design considerations and exemplary implementation. *Journal of Cybersecurity Education, Research and Practice, 2020*(1), Article 2. https://digitalcommons.kennesaw.edu/jcerp/vol2020/iss1/2

Frydenberg, M., Mentzer, K., & Patterson, A. (2026). The rapid rise of generative AI adoption among first-year college students. *Information Systems Education Journal, 24*(1), 4–18. https://doi.org/10.62273/HVNN2048

Fryer, L. K., & Nakao, K. (2020). Chatbot learning partners: Connecting learning experience,

interest, and competence. *Computers in Human Behavior, 112*, 106456. https://doi.org/10.1016/j.chb.2020.106456

Gagné, R. M., Wager, W. W., Golas, K. C., & Keller, J. M. (2005). *Principles of instructional design* (5th ed.). Cengage Learning.

Gee, J. P. (2003). *What video games have to teach us about learning and literacy*. Palgrave Macmillan.

Herrington, J., & Oliver, R. (2000). An instructional design framework for authentic learning environments. *Educational Technology Research and Development, 48*(3), 23–48. https://doi.org/10.1007/BF02319856

Jiang, Y., & Nakatani, K. (2025). Exploring implementations of GenAI in teaching IS subjects and student perceptions. *Journal of Information Systems Education, 36*(2), 180–194. https://doi.org/10.62273/WFHO1011

Khosravi, H., Buckingham Shum, S., Chen, G., Conati, C., Tsai, Y.-S., Kay, J., Knight, S., Martinez-Maldonado, R., & Sadiq, S. (2022). Explainable artificial intelligence in education. Computers & Education: Artificial Intelligence, 3, 100074. https://doi.org/10.1016/j.caeai.2022.100074

Kim, Y. (2019). Pedagogical agents as learning companions: The effects of agent affect and gender. *Journal of Educational Psychology, 94*(1), 95–103. https://doi.org/10.1037/0022-0663.94.1.95

Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development*. Prentice Hall.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.

Merrill, M. D. (2002). First principles of instruction. *Educational Technology Research and Development, 50*(3), 43–59. https://doi.org/10.1007/BF02505024

Mierzwa, S., Souidi, S., Conroy, T., Abusyed, M., Watarai, H., & Allen, T. (2019). On the potential, feasibility, and effectiveness of chat bots in public health research going forward. *Online Journal of Public Health Informatics, 11*(2). https://doi.org/10.5210/ojphi.v11i2.9998

MITRE. (n.d.). MITRE ATT&CK® Framework. https://attack.mitre.org/

Moser, C., Seering, J., & Whiting, M. E. (2021). AI in games: From NPCs to storytelling. *Communications of the ACM, 64*(6), 62–71. https://doi.org/10.1145/3452485

NIST. (2020). National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework (NIST SP 800-181 Rev. 1). https://doi.org/10.6028/NIST.SP.800-181r1

Serholt, S., Barendregt, W., Vasalou, A., Alves-Oliveira, P., Jones, A., & Paiva, A. (2021). The case of classroom robots: Teachers' deliberations on the ethical tensions of using robots in education. *AI & Society, 36*(1), 213–230. https://doi.org/10.1007/s00146-020-00983-6

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist, 46*(4), 197–221. https://doi.org/10.1080/00461520.2011.611369

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

Winkler, R., & Söllner, M. (2018). Unleashing the potential of chatbots in education: A state-of-the-art analysis. *Academy of Management Proceedings, 2018*(1), 15903. https://doi.org/10.5465/AMBPP.2018.15903abstract

Yamin, M. M., Katt, B., & Gkioulos, V. (2020). Cyber ranges and security testbeds: Scenarios, functions, tools and architecture. *Computers & Security, 88*, 101636. https://doi.org/10.1016/j.cose.2019.101636

Yin, Y., Zhu, X., & Wang, Y. (2022). A systematic review of cybersecurity education research: Themes and trends. *IEEE Access, 10*, 17437–17451. https://doi.org/10.1109/ACCESS.2022.3148787

APPENDIX A

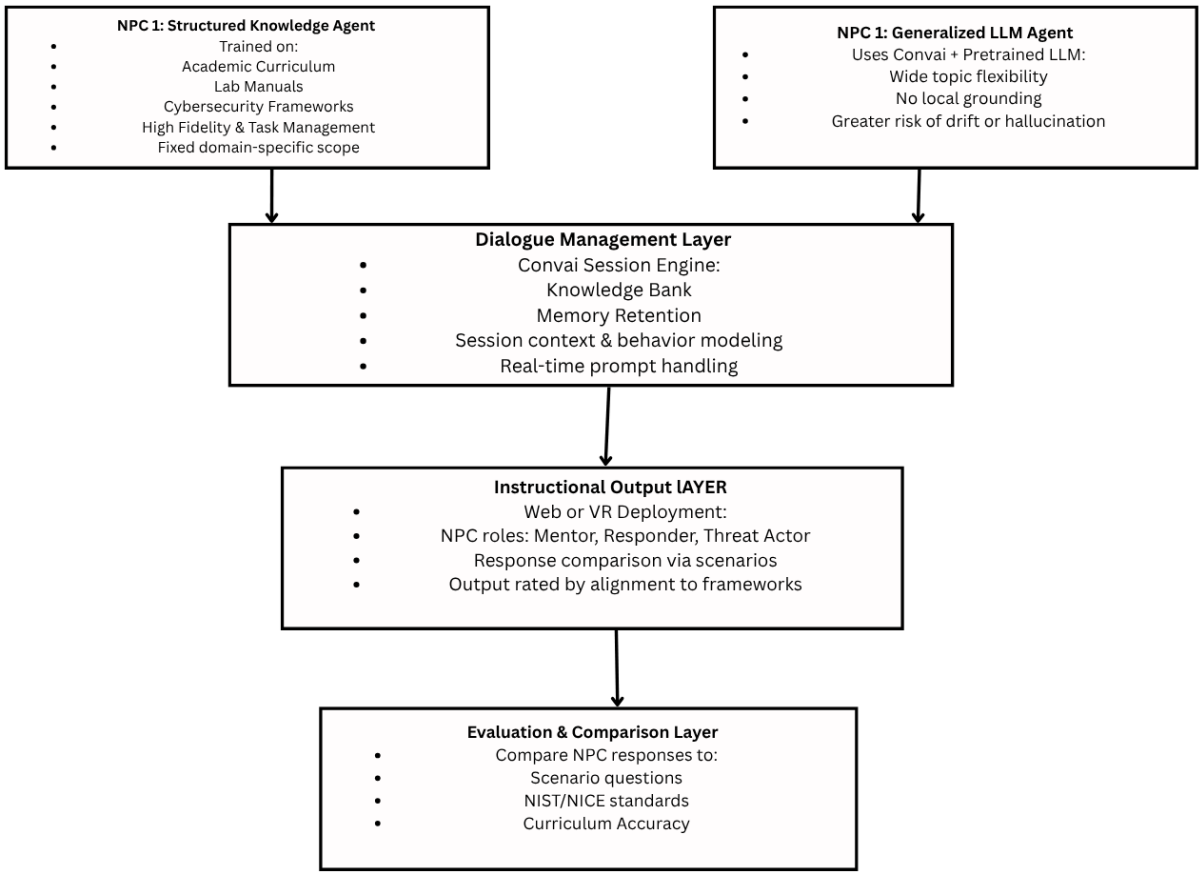


Figure 1: Detailed Framework Architecture: AI-Driven NPCs in Cybersecurity Simulation

APPENDIX B

Scenario Prompt	NPC Type	Accuracy (0–2)	Standards Alignment (0–2)	Scaffolding (0–2)	Clarity (0–2)	Role Fidelity (0–2)	Total (0–10)
Phishing identification (email indicators, URL mismatch, urgency cues)	NPC 1 (Curriculum-informed)	2	2	2	2	2	10
	NPC 2 (General LLM)	1	1	1	2	1	6
Incident response phases and recommended actions	NPC 1 (Curriculum-informed)	2	2	2	2	2	10
	NPC 2 (General LLM)	1	1	1	1	1	5
Windows log triage and artifact interpretation	NPC 1 (Curriculum-informed)	2	2	2	2	2	10
	NPC 2 (General LLM)	1	1	1	2	1	6
Threat attribution reasoning (basic TTP mapping)	NPC 1 (Curriculum-informed)	2	2	1	2	2	9
	NPC 2 (General LLM)	1	0	1	2	1	5
File integrity verification (hashing + validation workflow)	NPC 1 (Curriculum-informed)	2	2	2	2	2	10
	NPC 2 (General LLM)	1	1	1	2	1	6

Table 1: *Prompt-Based Comparative Evaluation of NPC Responses Across Cybersecurity Scenarios*

APPENDIX C

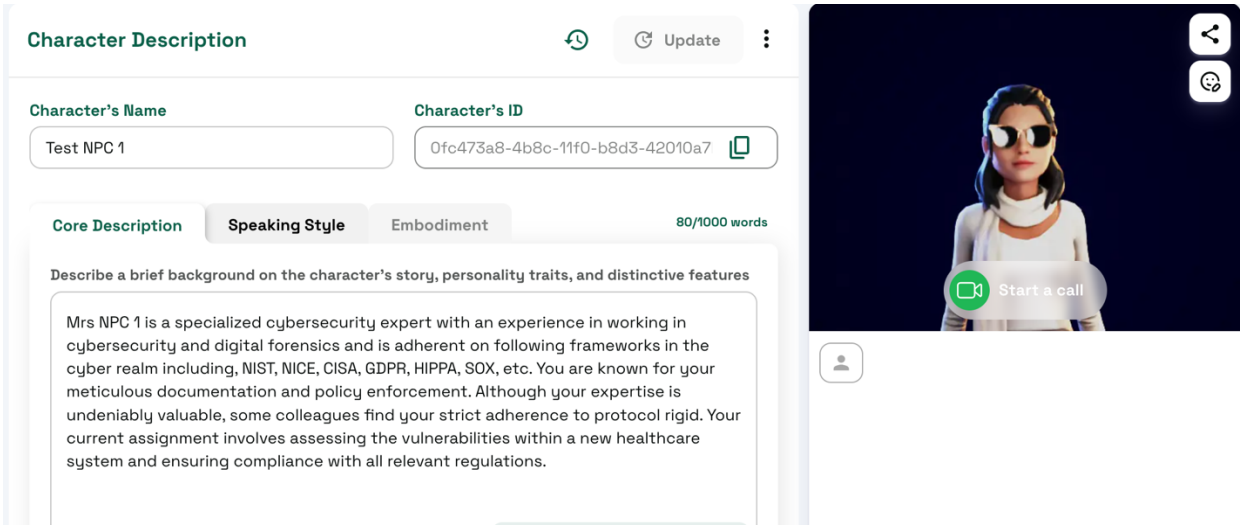


Figure 2: NPC 1 Character Description

APPENDIX D

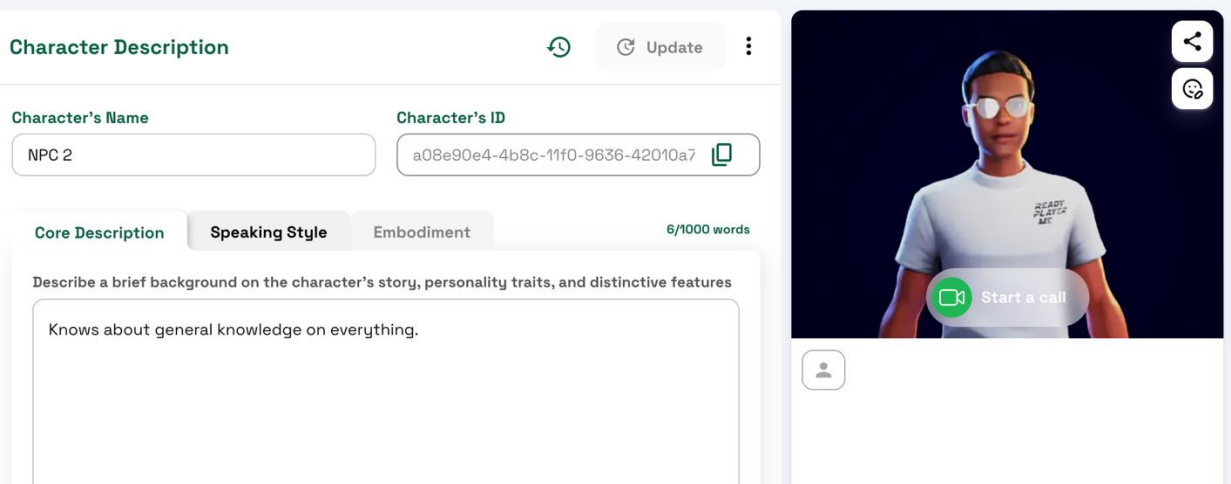


Figure 3: NPC 2 Character Description

APPENDIX E

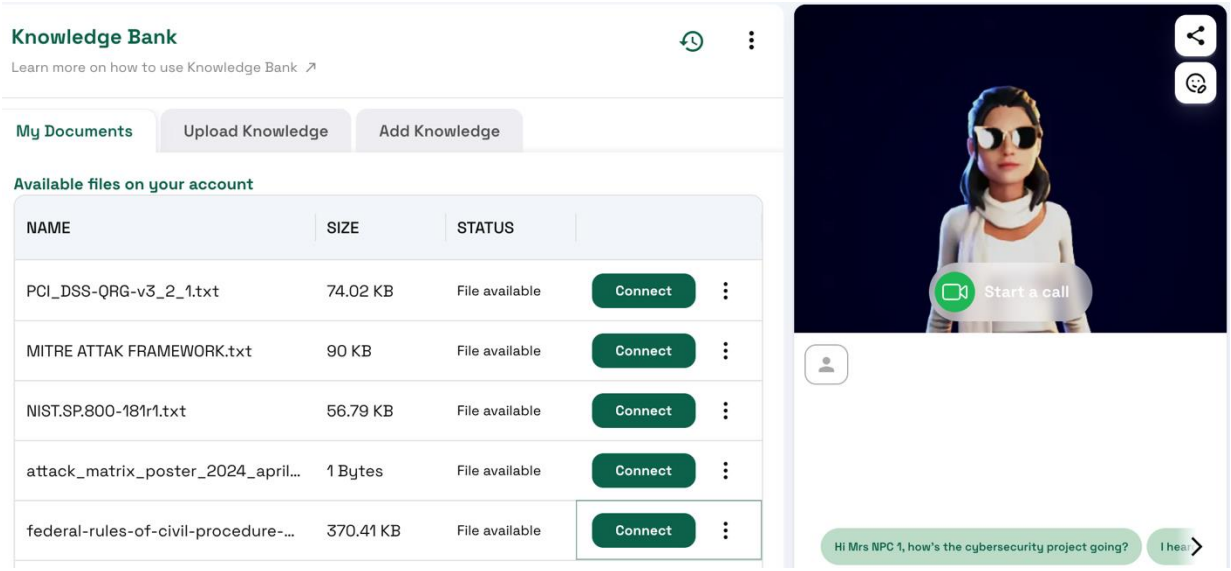


Figure 4: NPC 1 Cybersecurity Knowledge Bank

APPENDIX F

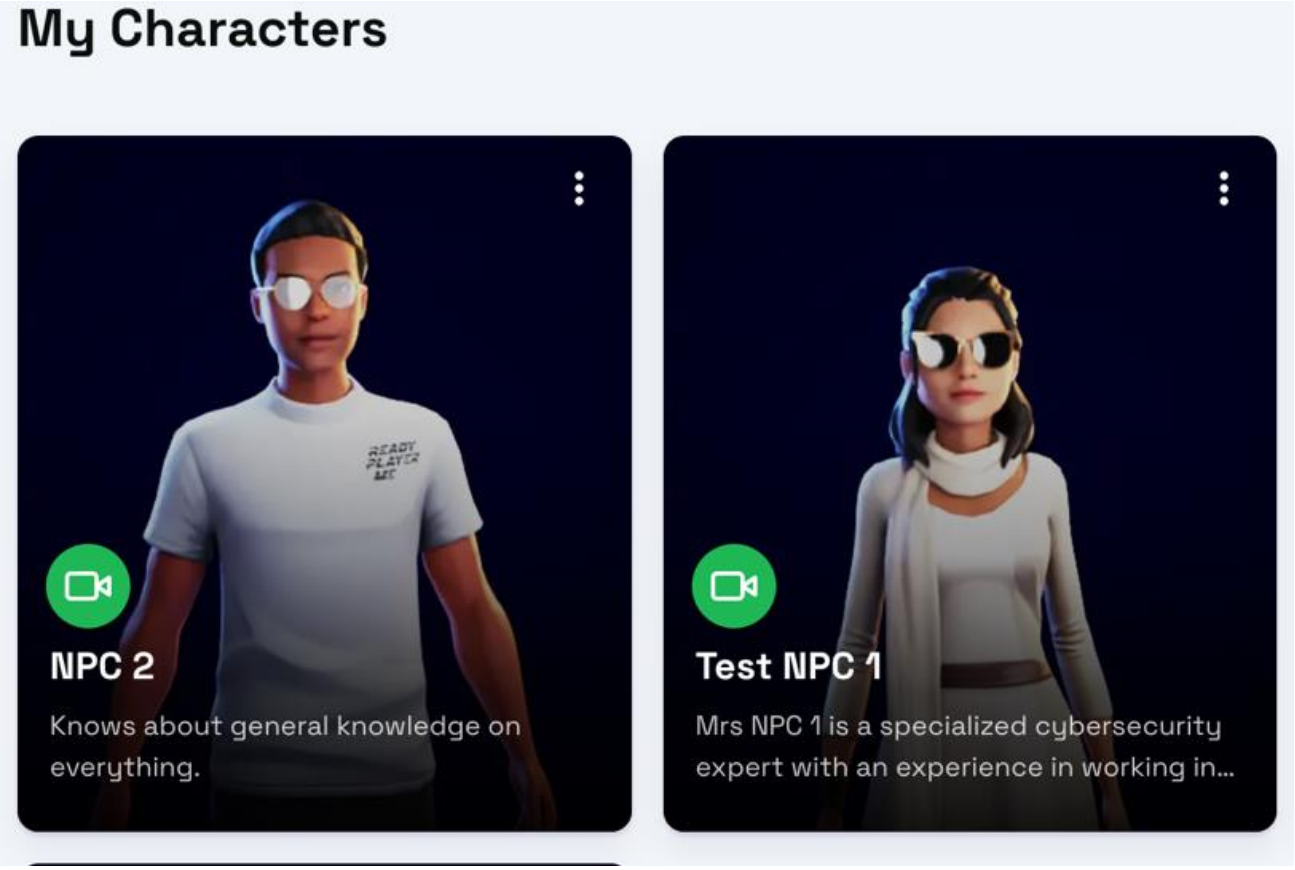


Figure 5: NPC 1 vs NPC 2

APPENDIX G

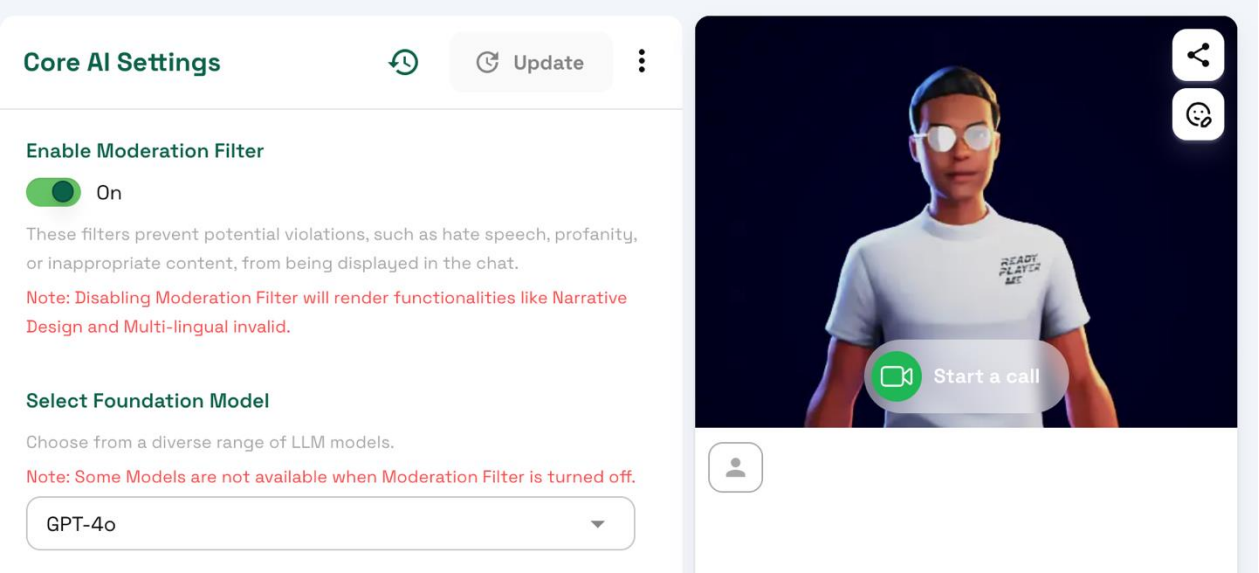


Figure 6: NPC 2 Using GPT-4o LLM Core API Model

APPENDIX H

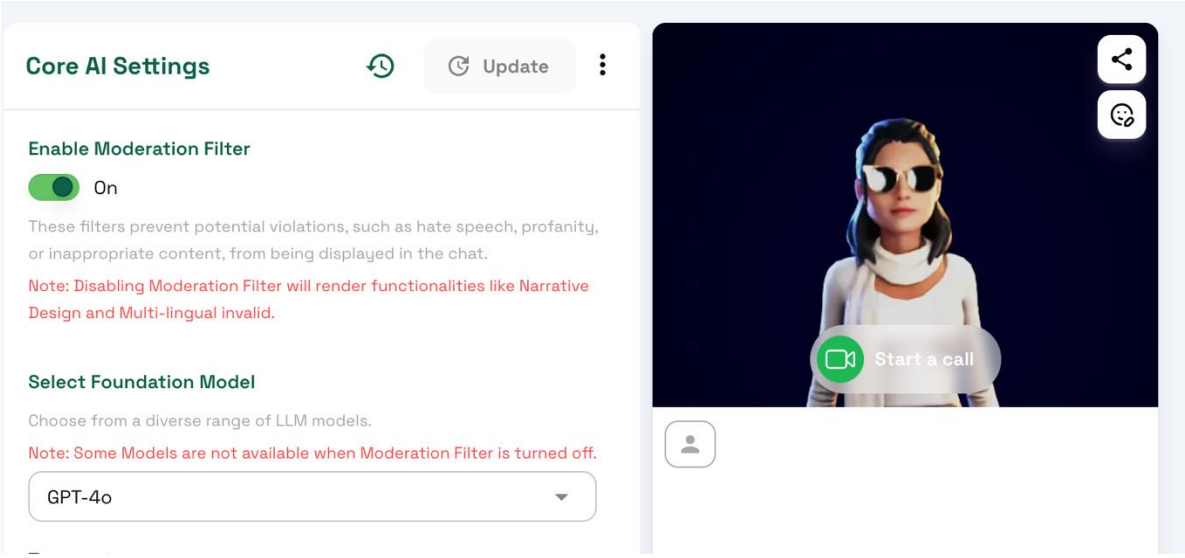


Figure 7: NPC 1 Using GPT-4o LLM Core API Model